

(NASA-TM-89418) RANGE FILTERING FOR
SEQUENTIAL GPS RECEIVERS WITH EXTERNAL
SENSOR AUGMENTATION (NASA) 22 p CSCI 17G

N87-20257

Unclas
63/04 45370

Range Filtering for Sequential GPS Receivers with External Sensor Augmentation

Russell Paielli

April 1987

Range Filtering for Sequential GPS Receivers with External Sensor Augmentation

Russell Paielli, Ames Research Center, Moffett Field, California

April 1987



National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

SUMMARY

The filtering of the satellite range and range-rate measurements from single channel sequential global positioning system (GPS) receivers is usually done with an extended Kalman filter which has state variables defined in terms of an orthogonal navigation reference frame. An attractive suboptimal alternative is range-domain filtering, in which the individual satellite measurements are filtered separately *before* they are combined for the navigation solution. The main advantages of range-domain filtering are decreased processing and storage requirements, and simplified tuning. Several range filter mechanization alternatives are presented, along with an innovative approach for combining the filtered range-domain quantities to determine the navigation state estimate. In addition, a method is outlined for incorporating measurements from auxiliary sensors into the range filtering scheme which is similar to the one used for incorporating the satellite measurements themselves. A method is also described for incorporating inertial measurements into the range filtering scheme as a process driver.

INTRODUCTION

The NAVSTAR Global Positioning System (GPS) will provide satellite-based, worldwide radio-navigation performance capabilities which will substantially exceed all other currently available systems. The potential applications include navigation for almost every type of land vehicle, surface ship, and aircraft for both military and civilian use.

The main development in this paper is a technique, referred to as range-domain filtering (or simply range filtering), for efficient, suboptimal Kalman filtering of measurements from single-channel (or dual-channel) sequential GPS receivers. These techniques are potentially applicable to all navigation applications in which a sequential GPS receiver is used, in either conventional or differential mode. The uses may eventually include low-cost civilian applications such as automotive, marine, general aviation, or helicopter navigation; or low-volume, low-power military applications such as tactical missiles, guided bombs, and hand-carried navigation units. With the satellite dwell times for sequential receivers expected to drop to one-eighth of a second or less, these receivers will become even more attractive for aerospace and automotive applications. In addition, the integration of low-cost sequential GPS receivers with low-cost inertial navigation systems may eventually result in economical navigation systems with good dynamic *and* long-term performance.

The intent of this paper is to present a basic analytical foundation for range-domain filtering. A brief outline of the conventional navigation-domain filtering (or simply navigation filtering) approach is presented for background, and theoretical considerations regarding the suboptimality of range filtering are discussed. Four different range filter mechanization alternatives are presented in addition to an efficient new technique for converting the filtered range-domain estimates to the navigation-domain state estimate. Several attractive

methods for incorporating external sensors, such as altimeters, into the range filtering scheme are discussed. Finally, a method is given for integrating an inertial navigation system (INS) into the range filtering scheme. The performance of the proposed concepts was not tested in the present paper, but the performance of range filtering for stand-alone GPS was tested by simulation in a related paper, reference 1, by the same author.

BACKGROUND

Currently the most common method by far of filtering single-channel sequential GPS receiver measurements for dynamic applications is what will be referred to as conventional navigation-domain filtering (or simply navigation filtering). Usually, some form of kinematic Kalman tracking filter is used, based on a linear state equation and a nonlinear measurement equation. The filter state is defined in terms of some orthogonal navigation reference frame such as a locally level frame.

Typically, the state vector consists of three position components (x, y, z) , user clock bias (b) , three velocity components $(\dot{x}, \dot{y}, \dot{z})$, and user clock frequency offset (\dot{b}) , for a total of eight state variables; sometimes three acceleration components $(\ddot{x}, \ddot{y}, \ddot{z})$ are also modeled, for a total of eleven state variables. The nonlinear measurement equation is linearized about the current state estimate to form an extended Kalman filter with raw measurement perturbations treated as linear combinations of the state-variable perturbations. The observation matrix consists of a time-varying Jacobian matrix, often referred to as the direction cosine matrix, which describes the satellite geometry and velocities. (Note that since multiplexed or multichannel receivers capable of continuous carrier phase-tracking are not considered, the deltarange measurements will be treated as discrete velocity measurements which are sampled simultaneously to the pseudorange measurements.)

An alternative to the conventional navigation-domain filtering is range-domain filtering, which is introduced by Van Graas in reference 2. As for navigation filtering, a kinematic Kalman tracking filter is used, but now a separate filter is required for each satellite in use. The state vector consists of the pseudorange and its time derivatives; typically the first one or two derivatives will be used, giving two or three state variables, respectively. The filtered range-domain quantities are then combined for the navigation solution.

Figure 1 illustrates the overall arrangement of the proposed navigation system. The satellite-sequencing mechanism, which involves pseudorandom noise codes and sophisticated hardware and software, is shown simply as a rotary switch accessing the satellites sequentially. Although four separate range filters are shown for clarity, the actual filter software may be common for all of them. The range-filter selection switch, which operates in synchronization with the satellite sequencing switch, then corresponds to the software function of loading the state vector and noise-covariance matrices for the processing of the current satellite measurements by the filter.

The four (or more) two- or three-variable range filters perform basically the same function as do the eight- or eleven-variable navigation filters, respectively. Since the processing and storage requirements of a Kalman filter are roughly proportional to the cube and square, respectively, of the state size, the advantage of range filtering in this respect is clear, even though more filters are required. Exact figures are difficult to obtain *a priori* concerning the number of flops executed by, and storage cells required by, each type of filter. This is true especially for the range filters because of the additional computation necessary for converting the filtered ranges to the navigation solution, as well as other overhead such as filter initialization at each new

satellite acquisition. Incidentally, the savings may be traded back for improved performance by implementing some adaptive tuning algorithm, which is more likely to be practical for the low-order range filters than it is for the navigation filters because of the larger number of process-noise covariance elements in the latter.

Suboptimal Filtering

A conceivable advantage of the navigation filters over the range filters is that they are potentially capable of taking into account the cross-correlations of measurement noise from one satellite to another if they are significant and known. However, since the satellite measurement updates are sequential, this capability no longer exists. Even if it did, there are other drawbacks. Since the cross-correlations depend on satellite geometry, they vary with time and hence require additional computation. Furthermore, once the estimated bias levels of the measurement noise have been subtracted out using propagation-delay models or differential corrections, as described in reference 3, the remaining noise component is virtually uncorrelated from one satellite to another for any reasonable satellite geometry. The capability of the navigation filter to utilize measurement-noise cross-correlation between satellites is therefore of negligible value.

For stand-alone GPS the treatment of process noise (unmodeled vehicle dynamics) is somewhat arbitrary because it consists predominantly of noise only during steady-state flight, is due, for example, to gust effects; during maneuvers the process noise is mainly a function of the control inputs, with a comparatively small additional noise component. Since the control inputs are unpredictable as far as the filter is concerned, the process noise is treated as white noise simply for mathematical expedience. Nevertheless, suppose that in some probabilistic sense it may be considered noise. A set of principal axes could then be found, in which each mode is statistically decoupled from the other modes; that is, the process-noise matrix could be diagonalized by some linear transformation. Intuitively, it would seem that the principal process-noise axes would most likely correspond to a ground-track reference frame. (The term "ground-track reference frame" will be used to represent a locally level frame which has its horizontal x-axis aligned with the vehicle ground-track.) In any other coordinate frame, cross-correlations would generally be assumed to be nonzero.

The range-domain filters may be thought of as navigation-domain filters which have been transformed to a nonorthogonal coordinate frame corresponding to the satellite directions, then separated into decoupled filters corresponding to each satellite, with the state-error covariance cross-correlation elements between satellites simply being discarded. Ideally, the cross-correlation elements would be zero; that is, the state-error covariance matrix, *after* transformation to the satellite frame and *before* decoupling, would be block-diagonal, where the blocks correspond to the individual satellites. Since the satellite coordinate frame is not a principal frame, however, the discarded cross-correlation elements are not generally zero.

The loss of the process-noise cross-correlation information, which was sacrificed for decreased complexity, is the reason that range filtering is a suboptimal approach. However, since the true process noise for unaided GPS is not really white noise, as is assumed by the Kalman filter, it may turn out that the sacrifice in optimality is not substantial or perhaps not even significant. Interestingly, if the process noise would be modeled as spherical anyway, that is, if the process-noise covariance value for each of the navigation axes is identical, then range filtering is almost equivalent to navigation filtering because then the principal axes are arbitrary. The only difference in that case is due to the fact that the satellite coordinate frame is nonorthogonal, a fact which becomes more significant the worse the dilution-of-precision (DOP) terms become. Note that the clock-error process-noise terms decompose ideally into the range domain because they

represent, by definition, errors directly along the satellite lines-of-sight. Thus, the degree to which range filtering is suboptimal depends on the degree to which the process noise is nonspherical and on the satellite DOP.

The criterion for decoupling the filter into separate filters for each satellite, then, is how nearly block-diagonal the full range-domain *state-error* covariance is. The *state-error* covariance depends, of course, on both the *process-noise* covariance and the *measurement-noise* covariance. Since the satellite coordinate frame is, in some sense, a principal frame for the *measurement-noise* covariance, the measurements themselves tends to "draw" the satellite frame toward the condition of being a principal frame with a block-diagonal *state-error* covariance. Thus, even though range filtering is not rigorously optimal, it is theoretically "closer" to optimal than if the conventional navigation filter were decoupled in some other arbitrary way, e.g., if it were decoupled by the navigation axes themselves.

FILTER MECHANIZATION

The range filters are based on the usual Kalman filter dynamics and measurement models for linear, time-invariant systems:

$$\mathbf{x}(k+1) = \Phi \mathbf{x}(k) + \mathbf{w}(k) \quad (1)$$

$$\mathbf{z}(k) = \mathbf{H}\mathbf{x}(k) + \mathbf{v}(k) \quad (2)$$

where \mathbf{x} is the state vector, Φ is the state-transition matrix for one discrete time-step, \mathbf{w} is the process-noise vector, \mathbf{z} is the measurement vector, \mathbf{H} is the observation matrix, \mathbf{v} is the measurement-noise vector, and k is the discrete time index. The noise statistics are described by

$$E[\mathbf{w}] = \bar{\mathbf{w}} \quad (3)$$

$$E[\mathbf{v}] = \bar{\mathbf{v}} \quad (4)$$

$$E[\mathbf{w}(k)\mathbf{w}(j)^t] = \mathbf{Q}(k)\delta_{kj} \quad (5)$$

$$E[\mathbf{v}(k)\mathbf{v}(j)^t] = \mathbf{R}(k)\delta_{kj} \quad (6)$$

$$E[\mathbf{w}(k)\mathbf{v}(j)^t] = \mathbf{0} \quad (7)$$

that is, the process noise \mathbf{w} and the measurement noise \mathbf{v} are statistically independent white noises with covariances \mathbf{Q} and \mathbf{R} , respectively.

Although the specifications of the Kalman filter mechanizations which follow are independent of the particular algorithmic implementation of the filter, the standard Kalman filter equations are summarized for reference in the Appendix. While the square-root Kalman filter algorithms have better numerical stability than the standard Kalman filter algorithm, numerical characteristics were not of primary concern in this study. Also, double-precision (eight byte) arithmetic is required for the state equations in this Kalman filter application because small perturbations of large numbers are processed, but that does not apply for the covariance equations. Of course, since the same processor will most likely be used for the equations of both the state and the covariance, the covariance equations will also get double-precision processing. With double-precision arithmetic the standard Kalman filter is roughly equivalent numerically to the square-root formulations with single-precision arithmetic.

The measurement-noise bias \bar{v} for GPS consists of the bias component of the various error sources such as ionospheric and tropospheric propagation delay error, multipath error, and satellite ephemeris, clock and group delay errors. These are subtracted out as well as possible by using differential corrections or by using propagation delay estimates based on the satellite data message. The process-noise bias \bar{w} is assumed to be zero for stand-alone GPS.

Although many different range-filter mechanizations are possible, the four considered to be potentially attractive are summarized below. The nomenclature is defined in conjunction with equations (1) and (2) above, with additional definitions as follows: r is the "true" pseudorange and \dot{r} is the "true" deltarange. For convenience, the true pseudorange is defined as the true range uncorrupted by any measurement error with the exception of user-clock bias and hardware bias; similarly, the true deltarange is defined as the true range rate corrupted only by the user-clock frequency offset. The measured pseudorange is denoted by r_m , Q is the process-noise covariance, and T is the time-interval, Δt , since the last update.

In each case, filter initialization is by default to some arbitrary "typical" state values with a large initial state-error covariance. At satellite switchover the filter reinitialization is more accurately performed based on the current navigation state estimate and the known satellite positions and velocities.

White-Noise-Jerk Filter

The white-noise-jerk (WNJ) filter treats the range jerk as white noise. The fundamental assumption is that

$$E[j(t)j(t+\tau)] = \sigma_j^2 \delta(\tau)$$

where j is the range jerk (third time-derivative of range), σ_j^2 is the jerk covariance and δ is the Dirac delta generalized function. Referring to equations (1) and (2), the filter is defined by

$$\begin{aligned} \mathbf{x} &\equiv \begin{bmatrix} r & \dot{r} & \ddot{r} \end{bmatrix}^t \\ \Phi &\equiv \begin{bmatrix} 1 & T & \frac{1}{2}T^2 \\ 0 & 1 & T \\ 0 & 0 & 1 \end{bmatrix} \\ \mathbf{z} &\equiv \begin{bmatrix} r_m & \dot{r}_m \end{bmatrix}^t \\ \mathbf{H} &\equiv \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \\ \mathbf{Q} &= \sigma_j^2 \begin{bmatrix} \frac{1}{5}T^6 & \frac{1}{4}T^5 & \frac{1}{3}T^4 \\ \frac{1}{4}T^5 & \frac{1}{3}T^4 & \frac{1}{2}T^3 \\ \frac{1}{3}T^4 & \frac{1}{2}T^3 & T^2 \end{bmatrix} \end{aligned}$$

Exponentially Correlated Acceleration Filter

The exponentially correlated acceleration (ECA) filter, developed by Singer in reference 4 for radar tracking of manned, maneuvering aircraft targets, treats the range acceleration as an exponentially correlated

random process. Applying this concept to the scalar range acceleration results in the fundamental assumption that

$$E[\ddot{r}(t)\ddot{r}(t+\tau)] = \sigma_{\ddot{r}}^2 e^{-\alpha|\tau|}$$

where $\sigma_{\ddot{r}}^2$ is the range-acceleration covariance and $1/\alpha$ is the acceleration time constant. The ECA filter is a generalization of the WNJ filter, which it reduces to when α approaches infinity.

By writing the continuous state equations, augmenting them by a first-order whitening filter, then discretizing the equations, Singer shows that, referring again to equations (1) and (2), the filter is defined by

$$\begin{aligned} \mathbf{x} &\equiv \begin{bmatrix} r & \dot{r} & \ddot{r} \end{bmatrix}^t \\ \Phi &\equiv \begin{bmatrix} 1 & T & \frac{1}{\alpha^2}[\alpha T + e^{-\alpha T} - 1] \\ 0 & 1 & \frac{1}{\alpha}[1 - e^{-\alpha T}] \\ 0 & 0 & e^{-\alpha T} \end{bmatrix} \\ \mathbf{z} &\equiv \begin{bmatrix} r_m & \dot{r}_m \end{bmatrix}^t \\ \mathbf{H} &\equiv \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \\ \mathbf{Q} &= \sigma_{\ddot{r}}^2 \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \end{aligned}$$

where

$$\begin{aligned} q_{11} &= \frac{1}{\alpha^4} \left[1 - e^{-2\alpha T} + 2\alpha T + \frac{2}{3}\alpha^3 T^3 - 2\alpha^2 T^2 - 4\alpha T e^{-\alpha T} \right] \\ q_{12} &= \frac{1}{\alpha^3} \left[e^{-2\alpha T} + 1 - 2e^{-\alpha T} + 2\alpha T e^{-\alpha T} - 2\alpha T + \alpha^2 T^2 \right] \\ q_{13} &= \frac{1}{\alpha^2} \left[1 - e^{-2\alpha T} - 2\alpha T e^{-\alpha T} \right] \\ q_{22} &= \frac{1}{\alpha^2} \left[4e^{-\alpha T} - 3 - e^{-2\alpha T} + 2\alpha T \right] \\ q_{23} &= \frac{1}{\alpha} \left[e^{-2\alpha T} + 1 - 2e^{-\alpha T} \right] \\ q_{33} &= \left[1 - e^{-2\alpha T} \right] \end{aligned}$$

Although the above expressions are not trivial, they are easily implemented because they need be computed only once, off-line.

White Noise Acceleration Filter

The white-noise-acceleration (WNA) filter treats the range acceleration as a white noise. The WNA filter is very similar to the WNJ filter, except that it does not model a range-acceleration state variable. The fundamental assumption in this case is that

$$E[\ddot{r}(t)\ddot{r}(t+\tau)] = \sigma_{\ddot{r}}^2 \delta(\tau)$$

Referring again to equations (1) and (2), the filter is defined by

$$\begin{aligned} \mathbf{x} &\equiv \begin{bmatrix} r & \dot{r} \end{bmatrix}^t \\ \Phi &\equiv \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \\ \mathbf{z} &\equiv \begin{bmatrix} r_m & \dot{r}_m \end{bmatrix}^t \\ \mathbf{H} &\equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ \mathbf{Q} &= \sigma_{\dot{r}}^2 \begin{bmatrix} \frac{1}{3}T^4 & \frac{1}{2}T^3 \\ \frac{1}{2}T^3 & T^2 \end{bmatrix} \end{aligned}$$

Exponentially Correlated Velocity Filter

The exponentially correlated velocity (ECV) filter models the range rate as an exponentially correlated random process. The ECV filter is very similar to the ECA filter, except that it does not model a range-acceleration state variable. The underlying assumption in this case is that

$$E[\dot{r}(t)\dot{r}(t+r)] = \sigma_{\dot{r}}^2 e^{-\alpha|r|}$$

where $\sigma_{\dot{r}}^2$ is the variance of the range rate and $1/\alpha$ is the range-rate time constant.

Again by writing the continuous state equations, then discretizing them, referring again to equations (1) and (2), the resulting filter is defined by

$$\begin{aligned} \mathbf{x} &\equiv \begin{bmatrix} r & \dot{r} \end{bmatrix}^t \\ \Phi &\equiv \begin{bmatrix} 1 & \frac{1}{\alpha}[1 - e^{-\alpha T}] \\ 0 & e^{-\alpha T} \end{bmatrix} \\ \mathbf{z} &\equiv \begin{bmatrix} r_m & \dot{r}_m \end{bmatrix}^t \\ \mathbf{H} &\equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ \mathbf{Q} &= \sigma_{\dot{r}}^2 \begin{bmatrix} q_{11} & q_{12} \\ q_{12} & q_{22} \end{bmatrix} \end{aligned}$$

where

$$\begin{aligned} q_{11} &= \frac{1}{\alpha^2} [4e^{-\alpha T} - 3 - e^{-2\alpha T} + 2\alpha T] \\ q_{12} &= \frac{1}{\alpha} [e^{-2\alpha T} + 1 - 2e^{-\alpha T}] \\ q_{22} &= [1 - e^{-2\alpha T}] \end{aligned}$$

For each filter the given observation matrix \mathbf{H} is appropriate if both pseudorange and deltarange measurements are available; if only the former is available, the bottom row of \mathbf{H} would simply be deleted.

The process-noise matrix for all of the above filters is specified by the selection of only one or two parameters: σ , or σ and α . The selection of the process-noise correlation time $1/\alpha$ is based on judgment or optimization. Singer's guidelines for the choice of $1/\alpha$ are: 60 sec for a lazy turn, 20 sec for an evasive maneuver, and 1 sec for atmospheric turbulence. The selection of the process-noise scalar multiplier σ may be performed as follows. A process-noise ellipsoid may be defined around the vehicle with the principal axes aligned with the ground-track reference frame. The magnitude of the principal ellipsoid axes are chosen based on the expected process noise along each particular axis, or by an off-line parameter optimization technique. The process-noise scalar multiplier for each satellite is then determined on-line by projecting the satellite line of sight through the ellipsoid. This allows for differing process-noise magnitudes in each axis, with a smooth transition from one axis to another. For a real-time system, the geometric computations involved may not be worthwhile, but some simpler method with unsmooth transitions would be relatively easy to generate, or the process noise could simply be considered constant for all satellites, which would correspond to the special case of a spherical process-noise ellipsoid.

As for which filter mechanization is appropriate for which application, the following general considerations apply. For applications with slowly varying but substantial vehicle acceleration, the higher-order filters which model a range-acceleration state variable, will probably be most effective. For applications with rapidly varying acceleration, the lower-order filters which estimate only the directly observable range and range rate, will probably outperform the higher-order filters because they will not mismodel the acceleration. If the vehicle acceleration is small, the lower-order filters will probably perform just as well as the higher-order filters will, but with less computation and storage.

CONVERSION ALGORITHM

The question arises as to how to optimally incorporate the filtered range-domain estimates for each satellite as they become available. The estimates from each satellite in use may of course be extrapolated to the current time and used in an algebraic navigation solution such as that described in reference 5, but that would require discarding any process noise (unmodeled vehicle dynamics) occurring during the extrapolation period. The conventional navigation filters do not have this difficulty; they are capable of incorporating each individual satellite measurement set in a theoretically optimal fashion. The navigation filters actually perform two functions: they filter out noise *and* they perform the nonlinear range-to-navigation conversion. Thus a conventional navigation filter may be simplified to perform only the nonlinear conversion on the previously filtered range-domain quantities.

This will be shown for the eight-variable conventional navigation filter; however, if a position estimate is all that is needed, it may just as easily be obtained with a four-variable (x, y, z, b) navigation filter though with slightly reduced performance and no significant advantages. The resulting algorithm will be referred to as the conversion algorithm. The modification is to simply omit all of the usual recursive state-error covariance computations, i.e., the propagation of the Riccati equation, and to compute only the linear state-propagation equation and the nonlinear measurement-update equation,

$$\hat{\mathbf{x}}^-(k+1) = \Phi \hat{\mathbf{x}}(k) + \bar{\mathbf{w}}(k) \quad (8)$$

$$\hat{\mathbf{x}}(k) = \hat{\mathbf{x}}^-(k) + \mathbf{K}_i(k)[\mathbf{z}_i(k) - \bar{\mathbf{z}}_i(k)] \quad (9)$$

where the symbols are defined as follows: $\hat{\mathbf{x}}^-$ and $\hat{\mathbf{x}}$ are the filtered estimates of the navigation-domain state vector \mathbf{x} just before, and just after, the measurement update, respectively, with

$$\mathbf{x} \equiv \begin{bmatrix} x & y & z & b & \dot{x} & \dot{y} & \dot{z} & \dot{b} \end{bmatrix}^t \quad (10)$$

\mathbf{z}_i is the previously filtered state vector from the i -th satellite range filter (or rather, the first two elements of it, r and \dot{r}), used now as an equivalent "measurement input" to the navigation filter, and $\bar{\mathbf{z}}_i$ is the *predicted* value of \mathbf{z}_i based on the time-updated navigation state and the known satellite position and velocity,

$$\bar{\mathbf{z}}_i \equiv \begin{bmatrix} \bar{r}_i & \dot{\bar{r}}_i \end{bmatrix}^t \quad (11)$$

with the nonlinear prediction functions given by

$$\begin{aligned} \bar{r}_i &= [(x - x_{s,i})^2 + (y - y_{s,i})^2 + (z - z_{s,i})^2]^{1/2} + b \\ \dot{\bar{r}}_i &= \frac{(x - x_{s,i})(\dot{x} - \dot{x}_{s,i}) + (y - y_{s,i})(\dot{y} - \dot{y}_{s,i}) + (z - z_{s,i})(\dot{z} - \dot{z}_{s,i})}{[(x - x_{s,i})^2 + (y - y_{s,i})^2 + (z - z_{s,i})^2]^{1/2}} + \dot{b} \end{aligned} \quad (12)$$

where $x_{s,i}$ is the x -component of the i -th satellite position and where $\dot{x}_{s,i}$ is the x -component of the i -th satellite velocity, etc., and where x , y , z and b are the elements of the vector in equation (10) after update by equation (8). Also, Φ is the navigation-domain state-transition matrix for one discrete time-step,

$$\Phi \equiv \begin{bmatrix} \mathbf{I}_{(4 \times 4)} & T\mathbf{I}_{(4 \times 4)} \\ \mathbf{0}_{(4 \times 4)} & \mathbf{I}_{(4 \times 4)} \end{bmatrix} \quad (13)$$

The gain \mathbf{K}_i for update by the i -th satellite is given by

$$\mathbf{K}_i(k) = \mathbf{P}(k)\mathbf{H}_i(k)^t \mathbf{R}_i(k)^{-1} \quad (14)$$

which is one of the standard Kalman gain formulas, but in this context the variables have a somewhat different meaning, as follows. \mathbf{H}_i is the time-varying observation Jacobian matrix for the i -th satellite, often referred to as the direction cosine matrix,

$$\mathbf{H}_i = \begin{bmatrix} \mathbf{h}_i (1 \times 4) & \mathbf{0} (1 \times 4) \\ \mathbf{0} (1 \times 4) & \mathbf{h}_i (1 \times 4) \end{bmatrix} \quad (15)$$

where

$$\mathbf{h}_i = \begin{bmatrix} \cos \theta_{x,i} & \cos \theta_{y,i} & \cos \theta_{z,i} & 1 \end{bmatrix} \quad (16)$$

and $\cos \theta_{x,i}$ is the cosine of the angle between the navigation frame x -axis and the i -th satellite line of sight,

$$\cos \theta_{x,i} = \frac{x_{s,i} - x}{[(x - x_{s,i})^2 + (y - y_{s,i})^2 + (z - z_{s,i})^2]^{1/2}} \quad (17)$$

with similar definitions of course for the y - and z -components. An approximation in equation (15) involves disregarding the first three elements in the second row, which are nonzero in general but are small; the maximum value these elements can reach (for a stationary user) is the inverse of the 12-hr GPS orbital period, approximately 1.5×10^{-4} /sec. The equivalent "measurement-error" covariance, \mathbf{R}_i , is the state-error covariance

matrix from the i -th satellite range filter (or rather, the upper left 2×2 portion of it). The equivalent state-error covariance matrix in the navigation domain, \mathbf{P} , is based on the state-error covariances of the individual range-domain filters. That is, \mathbf{P} is computed by mapping the state error covariances in the range domain to the navigation domain, a purely algebraic transformation of the recursively computed range-domain covariances. It is assumed that the satellite geometry changes slowly enough so that it may be approximated as constant during the satellite sequencing cycle. The transformation is given by

$$\mathbf{P}(k) = \mathbf{H}(k)^\dagger \mathbf{R}(k) \mathbf{H}(k)^{\dagger t} \quad (18)$$

where \mathbf{H} is the composite observation matrix for all the satellites in use, that is

$$\mathbf{H} \equiv \begin{bmatrix} \mathbf{H}_1 \\ \mathbf{H}_2 \\ \vdots \end{bmatrix} \quad (19)$$

and \mathbf{R} is a block diagonal matrix consisting of the individual range-domain filter state-error covariances (actually, the upper left 2×2 portions of them, as in \mathbf{R}_i), that is

$$\mathbf{R} \equiv \text{diag}\{\mathbf{R}_1, \mathbf{R}_2, \dots\} \quad (20)$$

where the satellite ordering is the same as the ordering in definition (19). Finally, \mathbf{H}^\dagger is the pseudoinverse of \mathbf{H} , defined as

$$\mathbf{H}^\dagger \equiv [\mathbf{H}^t \mathbf{H}]^{-1} \mathbf{H}^t \quad (21)$$

which reduces to \mathbf{H}^{-1} with exactly four satellites in use and becomes singular with less than four satellites. The pseudoinverse has been used as a generalization to illustrate the fact that more than four satellites may be utilized if they are visible. This would improve the satellite geometry somewhat (decreased GDOP), but would slow the effective sequencing rate for the entire satellite group. Whether or not this would be advantageous depends on the sample rate and the magnitude of the process noise; the milder the process noise and the higher the sample rate, the more likely that the redundant satellites will be useful. With high-rate inertial measurements used as a process driver, as discussed in reference 1, the process noise would be the measurement noise of the inertial instruments, hence the redundant satellite capability would be very attractive in that case.

Note that equation (21) appears to require the the inversion of an 8×8 matrix in real-time; however, the special structure of \mathbf{H} (cf. equation 15) allows the inversion to be restructured for a real-time implementation so as to actually only require a 4×4 matrix inversion, reducing the computation time by roughly a factor of eight. In addition, the structures of \mathbf{H} and \mathbf{H}^\dagger allows them to be stored in one-quarter of the space that the full matrices would require.

Also, the fact that the satellite geometry changes relatively slowly allows several of the computations indicated above to be performed less often than the basic GPS sample rate. For example, the matrices describing the satellite geometry, \mathbf{H}_i , \mathbf{H} , and \mathbf{H}^\dagger , need not be computed at each discrete time-step. The predicted measurement vector $\bar{\mathbf{z}}_i$ must be computed at each time-step, but the square root in equation (12) need not be computed each time. An alternative is to compute the square root only occasionally, then at the other times to compute the distance from user to satellite by computing perturbations on this value. The perturbations would be calculated by projecting the predicted position perturbations of *both the user and the satellite* onto the satellite line-of-sight, using the direction cosines that have already been computed for equation (17) in

the process.

Since double-precision square roots and matrix inversion are computationally expensive the processor load is significantly reduced by decreasing the execution rates for equations (15) and (21), and the square root in equation (12). Since the required update rates for these are all based on the satellite geometry, they could all be updated together at some background rate which is substantially less than the primary GPS update rate. A satisfactory background rate must be empirically determined.

The matrix multiplication in equation (18) appears to be computationally formidable. If it had to be fully multiplied at each discrete time-step the efficiency of the range-filtering scheme would become far worse than it is for conventional navigation filtering, thus defeating the whole purpose of range filtering. Since the matrices have sparse structures the computational load could be reduced by a factor of 16.

More importantly, however, careful consideration reveals that the multiplication in equation (18) may be simplified considerably. By combining equations (14) and (18) it is apparent that the gain matrix is given by

$$\mathbf{K}_i = \mathbf{H}^\dagger \mathbf{R} \mathbf{H}^\dagger \mathbf{H}_i \mathbf{R}_i^{-1} \quad (22)$$

When exactly four satellites are in use, the effect of post-multiplying by $\mathbf{R} \mathbf{H}^\dagger \mathbf{H}_i \mathbf{R}_i^{-1}$ is to simply select the two columns of \mathbf{H}^\dagger corresponding to the current satellite in the sequence, thus rendering the matrix multiplication both in equations (18) and (14) completely unnecessary. This method is similar to the scalar update navigation algorithm of reference 6.

When more than four satellites are in use, the gain matrix computation becomes somewhat more complicated but may still remain efficient. The use of redundant satellites corresponds to an overdetermined system of equations, and the solution is a least-squares fit to the data. The individual satellites may then be weighted in the conversion algorithm according to their estimated accuracies. The definition of the pseudoinverse given by equation (21) implies equal weighting of the satellites, but the relative weighting may be set arbitrarily by using a weighting matrix \mathbf{W} in the definition, that is

$$\mathbf{H}^\dagger \equiv [\mathbf{H}^t \mathbf{W} \mathbf{H}]^{-1} \mathbf{H}^t \mathbf{W} \quad (23)$$

Ideally \mathbf{W} would just be equal to \mathbf{R} , the composite state-error covariance matrix of the range filters. The matrix inversion can then no longer be reduced to a 4×4 inversion. However, if the weighting is only changed at the background geometry-update rate, then the inversion need only be executed at that same background rate. Furthermore, the matrix multiplication in equation (22) need only be executed at the background rate also. The conversion algorithm would then have a fixed gain matrix for each satellite throughout the entire geometry-update period. Fortunately, there is no practical need to change the satellite weighting at a faster rate than the background geometry-update rate, so the conversion algorithm remains very efficient even with more than four satellites in use.

Navigation Initialization

Initialization of the conversion filter requires careful consideration. As usual, an estimate of the initial navigation coordinates is required, along with an estimate of the initial error covariance. However, the

algorithm discussed above for computing the equivalent navigation state error covariance from the range filter state error covariances cannot be used until the satellite sequencing cycle has been completed one time because otherwise the computed covariance will have no relation to the covariance of the initial estimate.

For the initial measurement update by each of the satellites in use, the updated navigation state error covariance could be calculated as

$$\mathbf{P}(k+1) = [\mathbf{P}^{-1}(k) + \mathbf{H}_i(k)^t \mathbf{R}_i(k)^{-1} \mathbf{H}_i(k)]^{-1} \quad (24)$$

in place of equation (18). However, this is too computationally intensive and not worth the burden just to start the filter. An alternative approach would be to simply take the initial updates by each satellite and process them all together. Once the last satellite in the sequence is initially accessed the range-domain state estimates and error covariances from each satellite would be time-propagated to the current time, then the state estimates would be loaded into a composite measurement vector,

$$\mathbf{z} \equiv \begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \\ \vdots \end{bmatrix} \quad (25)$$

The equations previously given could then be used as is except that \mathbf{z}_i would be replaced by \mathbf{z} in equation (9), and equation (14) would become

$$\mathbf{K}(k) = \mathbf{H}(k)^{-1} \quad (26)$$

This approach has the advantage over the use of equation (24) that the computation of \mathbf{P} is restructurable to a 4×4 , rather than an 8×8 , matrix inversion. Additionally, the initial state estimate may be set by default to a crude guess, and an initial state error covariance is not required.

Clock Coasting

Three-dimensional aircraft navigation with GPS typically requires the use of four satellites, the "extra" satellite providing information for the determination of user clock error. If, for a short time, only three satellites are available, e.g., as a result of temporary line of sight blockage or radio-frequency interference, a method called clock coasting may be used. The concept is based on the assumption that the user clock error "dynamics" are milder than the vehicle dynamics, thus the conversion filter may be allowed to quit tracking the clock error for a short period and navigate with only three satellites. The length of time for which this may reasonably be done depends, of course, on the user clock frequency stability.

The clock coasting method requires the filter to be programmed for the assumption that the user clock frequency offset (bias rate) remains constant during the coasting period. The proper method for achieving this requires some careful consideration. First of all, the matrices in equations (19) and (20) should only be loaded with values appropriate for the satellites in use, not including the temporarily lost satellite. However, with only three satellites available, the matrix to be inverted in equation (21) becomes singular. The way around this problem, consistent with the fact we are no longer attempting to measure the clock error, is to simply eliminate the fourth and eighth columns of the matrix in equation (15) for each satellite, which results in a 6×6 instead of 8×8 \mathbf{P} matrix. Then after the 6×6 \mathbf{P} matrix is computed, we simply convert it back to an 8×8 matrix by adding a fourth row and column, and an eighth row and column, all consisting of zeros. This implies that the clock error is perfectly known and hence does not get updated by the measurement, which is

consistent with the clock coasting assumption.

This operation may be expressed mathematically as follows. For convenience, define the 8×6 matrix $\mathbf{I}_{8 \times 6}$ as an 8×8 identity matrix with the fourth and eighth columns removed, and define $\mathbf{I}_{6 \times 8} \equiv \mathbf{I}_{8 \times 6}^t$. The conversion filter equations all remain unchanged except equation (18), which becomes

$$\mathbf{P}(k) = \mathbf{I}_{8 \times 6} [\mathbf{H}(k) \mathbf{I}_{8 \times 6}]^t \mathbf{R}(k) [\mathbf{H}(k) \mathbf{I}_{8 \times 6}]^t \mathbf{I}_{6 \times 8} \quad (27)$$

Interestingly, if the user clock has adequate frequency stability it may be advantageous to use clock coasting even if four or more satellites are available. By using only three satellites, the effective sample rate for the whole group is higher, providing improved response in dynamic flight conditions. For a clock error update, the fourth satellite could be accessed intermittently, with the original state error covariance equations used during the full satellite sequencing cycle in which the fourth satellite is accessed. Presumably, the clock error should be updated when its predicted error covariance approaches that for the most accurate spatial axis, which could be computed on-line. During the sequencing cycles when only three satellites are accessed, the satellite selection criterion would be an optimized PDOP for the three satellites.

EXTERNAL SENSOR INTEGRATION

For many GPS applications it is desirable to augment the receiver with additional sensors such as inertial instruments and/or a radar- or baro-altimeter. Such measurements may be incorporated into the filter either as process inputs or as measurement inputs. The former affect the Kalman filter time-update (one-step prediction) and the latter affect the filter measurement update.

Auxiliary Measurement Inputs

For aircraft precision approach and landing the vertical axis has the most critical accuracy requirements. Vertical accuracy may be improved with the use of an altimeter or a vertical accelerometer attached to some vertical reference such as an artificial horizon. Measurements from auxiliary sensors of this type must be incorporated into the range filtering scheme as actual measurement inputs.

Auxiliary measurement inputs of this sort fit well in conventional navigation filtering schemes because they are easily converted to, if not already equivalent to, direct observations of filter states. Unfortunately, the same does not apply for range-domain filtering. The auxiliary measurements are not generally made along an axis coinciding with a satellite line of sight. Furthermore, the nonzero satellite velocities would complicate any attempt to project the auxiliary measurements onto the individual range filters. In order to circumvent these difficulties and to conveniently incorporate auxiliary measurements into the range filtering scheme, the concept of a virtual satellite is now introduced.

A virtual satellite is an imaginary satellite located at a convenient position, and with a convenient (zero) velocity. A virtual satellite would be defined for each auxiliary sensor. The convenient position is such that the satellite is at some arbitrary, large distance and its line of sight is coincident with the sensitive axis of the particular sensor. As an example, for an altimeter or vertical accelerometer, the satellite would always be directly above (or below) the vehicle. Notice that this substantially enhances the effective satellite geometry, especially improving the vertical dilution of precision (VDOP), which happens to be both the most critical

axis for aircraft landing and the least accurate axis for stand-alone GPS.

The auxiliary measurements may then be treated as equivalent to redundant satellite measurements: they may be filtered using appropriate "range filters", and then the output states of the filters may be combined with the *actual* satellite measurements using the range to navigation conversion technique outlined above. For distance-measuring sensors such as altimeters, the measured quantities would have to be subtracted from some arbitrary datum to obtain the equivalent satellite range measurement, but since the virtual satellites are defined as having zero velocity, measurements from velocity and acceleration sensors may be used as is. The fact that the auxiliary measurements are not affected by user clock bias, as are the actual satellite measurements, is accounted for in the observation matrix H_i (cf. equation 15), which will have a clock bias coefficient of zero instead of one. For velocity and acceleration type sensors, the equivalent range filters must of course be provided with appropriate initial conditions if they are to provide an equivalent range estimate. The initial conditions may be updated, i.e., the filter may be reset, at each GPS update (or less often if desired), based on the navigation solution and the virtual satellite position. Note also that the virtual "range filters" may operate at an arbitrarily higher sample rate than the GPS sample rate, if desired, and they may also be used to update the navigation state estimate at any desirable rate.

Note that each virtual satellite allows navigation with one less real satellite. For example, with the use of an altimeter only three satellites are required for three-dimensional navigation; however, accuracy will suffer if the altimeter is not as accurate as the satellites. With the use of clock coasting in addition to an altimeter, three-dimensional navigation may actually be accomplished for a short period of time with only two satellites available. The implications for reliability and redundancy are obvious.

Process Inputs

A promising application for GPS is to integrate it with a strapdown inertial navigation system (INS) and use it to periodically reset the accumulated INS errors. Since INS systems have excellent short term characteristics, low-cost sequential GPS receivers may be more than adequate for many GPS-INS applications. In this scheme the measurements from the INS can hardly be considered as auxiliary inputs to the GPS receiver, but the same conceptual framework applies. However, since the inertial measurements constitute twice-differentiated observations of the entire navigation "space", it becomes very attractive to incorporate them into the range filtering scheme as process inputs.

This would be implemented as follows. The navigation state estimate would be propagated from one satellite update to the next, using the acceleration measurement as the process input. This means that the average acceleration over the satellite dwell time, as computed by the strapdown system, would be used as a process-noise bias in the time-update of the eight-variable conversion algorithm, i.e., $\bar{w}(k)$ in equation (8) would become $\bar{a}(k, k+1)$, the average acceleration between times k and $k+1$. The expected pseudorange and deltarange measurements for the next satellite in the sequence would then be computed and used as the predicted state in the corresponding range filter, with the *process-noise* covariance for the range filter based on the acceleration *measurement-noise* covariance and the geometry. The state time-update is executed in the navigation domain and the measurement update is executed in the range domain. The process-noise covariance estimate for the range filter may then be considerably reduced, thereby tightening the filter bandwidth for increased noise rejection. The appropriate range filter mechanization for this approach is the previously defined white-noise-acceleration filter.

Several additional points are worth mentioning for completeness even though they are beyond the scope of the present study. First, the inertial measurements may be used to narrow the GPS receiver tracking-loop bandwidths, in addition to the filter bandwidths, for improved receiver accuracy. Secondly, an adaptive algorithm could be designed to perform in-flight alignment of the INS attitude reference system based on the range filter residuals.

Filter Aiding Inputs

Finally, one more possible use of external sensors is as an aid for the range filters. This method involves the use of an aircraft attitude-and-heading reference system (AHRS) or vehicle control-input sensors.

For example, if aircraft roll attitude information is available (as it is for nearly all general aviation aircraft), the assumption of coordinated turns, if reasonable, allows the flightpath to be predicted reasonably well. For a helicopter, the pitch attitude may also be used to predict longitudinal acceleration. Alternatively, signals from displacement transducers installed on the pilot's control devices may be fed to a state estimator programmed to predict the transient response of the aircraft.

With these filter aiding schemes the range-domain state variables for each satellite are more accurately predictable than with the simple constant-velocity or constant-acceleration models, thus allowing the use of narrower filter bandwidths for improved noise rejection.

SUMMARY AND CONCLUSIONS

An alternative to the conventional navigation filtering for sequential GPS receivers with or without auxiliary sensors has been developed. The technique involves decoupling the navigation filter into separate range-domain filters for each satellite. In decoupling the filters, some process-noise cross-correlation terms are implicitly discarded, but the actual nature of the process noise (for stand-alone GPS) may be such that the discarded information is insignificant. Measurement-noise cross-correlation information between satellites is not affected because it cannot be used with sequential measurement updates. The potential advantages of the range filtering scheme are decreased numerical processing and storage requirements, and simplified filter tuning. The latter implies that adaptive tuning may be more practical for range-domain filters than it is for navigation-domain filters, which is an important consideration.

Four different range filter mechanizations were presented: two of them estimate only the directly observable range and range-rate, the other two also model the range-acceleration. General guidelines for choosing the appropriate filter for particular applications were given.

A new technique is presented for converting the range-domain state estimates to the navigation state estimate. The conversion filter, as it is referred to, allows the filtered range-domain state estimates from each satellite to be incorporated sequentially into the navigation state estimate as they become available. The conversion algorithm consists basically of a conventional navigation filter but without the recursive state error covariance computations. Instead, an equivalent state error covariance matrix is computed based on the individual range-filter state-error covariances and on the satellite geometry. A conversion filter initialization

algorithm is given and a method is outlined to program the conversion filter for clock coasting so that navigation may continue for a short time if the number of accessible satellites becomes deficient.

Finally, it was pointed out that measurements from external sensors may be integrated into the filtering scheme in several ways, depending on the sensor type. By using a mathematical concept referred to as a virtual satellite, altimeter measurements, for example, may be treated equivalently to satellite range measurements. In an integrated GPS-INS system, on the other hand, the inertial measurements may be used as a navigation-domain process driver in the conversion filter, and the predicted navigation-domain state estimate may then be used to compute the predicted range-domain state estimates, thus allowing the use of narrower range-filter bandwidths for improved noise rejection. Filter-aiding schemes were also pointed out: aircraft attitude measurements may be used to infer the predicted flightpath if a coordinated turn assumption is reasonable; alternatively, direct measurements of vehicle control inputs may be fed into an aircraft-dynamics state estimator to predict the flightpath. The filter aiding schemes are limited in performance by both the measurement noise and the uncertainty in the flight path prediction models, but they are likely to be substantially more accurate than the simple kinematic models which must otherwise be used, thus allowing the range filter bandwidths to be narrowed for improved noise rejection.

REFERENCES

1. Paielli, R. A.: Range Filtering for Sequential GPS Receivers. Proc. of the Institute of Navigation National Technical Meeting, Jan. 1987.
2. Van Graas, F.: Discrete Filtering Techniques Applied to Sequential GPS Range Measurements. NASA Contract NAS2-11969, TM H-6, Jan. 1985.
3. Edwards, F. G.; and Loomis, P. V. W.: Civil Helicopter Flight Operations Using Differential GPS. J. of the Inst. of Navigation, vol. 32, no. 4, Fall 1985.
4. Singer, R. A.: Estimating Optimal Tracking Filter Performance for Manned Maneuvering Targets. IEEE Trans. on Aerospace and Electronic Systems, vol AES-6, no. 4, July 1970.
5. Bancroft, S.: An Algebraic Solution of the GPS Equations. IEEE Trans. on Aerospace and Electronic Systems, vol. AES-21, no. 7, Jan. 1985.
6. Wong, G. A.; and Joglekar, A. N.: Scalar Update Navigation Algorithm for a Low-Cost NAVSTAR GPS Receiver. IEEE, Proc. of the Third Digital Avionics Systems Conf., pp. 64-69, Fort Worth, Texas, Nov. 1979.
7. Gelb, A., Ed.: *Applied Optimal Estimation*. The M.I.T. Press, Cambridge, Massachusetts, 1974.
8. Brown, R. G.: *Introduction to Random Signal Analysis and Kalman Filtering*. John Wiley and Sons, New York, 1983.
9. Bryson, A. E.; and Ho, Y.: *Applied Optimal Control*, Hemisphere Publishing Corp., 1975.

10. Kronhamn, T. R.: Geometric Illustration of the Kalman Filter Gain and Covariance Update Algorithms. IEEE Control Systems Magazine, May 1985.

APPENDIX

STANDARD KALMAN FILTER

Consider the linear system with additive Gaussian noise given by

$$\begin{aligned}\mathbf{x}(k+1) &= \Phi\mathbf{x}(k) + \mathbf{w}(k) \\ \mathbf{z}(k) &= \mathbf{H}\mathbf{x}(k) + \mathbf{v}(k)\end{aligned}$$

where \mathbf{x} is the state vector, Φ is the state transition matrix for one discrete time step, \mathbf{w} is the process-noise vector, \mathbf{z} is the measurement vector, \mathbf{H} is the observation matrix, \mathbf{v} is the measurement-noise vector, and k is the discrete time index.

Suppose that the noise statistics are described by

$$\begin{aligned}E[\mathbf{w}] &= \bar{\mathbf{w}} \\ E[\mathbf{v}] &= \bar{\mathbf{v}} \\ E[\mathbf{w}(k)\mathbf{w}(j)^t] &= \mathbf{Q}(k)\delta_{kj} \\ E[\mathbf{v}(k)\mathbf{v}(j)^t] &= \mathbf{R}(k)\delta_{kj} \\ E[\mathbf{w}(k)\mathbf{v}(j)^t] &= 0\end{aligned}$$

that is, the process noise \mathbf{w} and the measurement noise \mathbf{v} are statistically independent white noises with covariances \mathbf{Q} and \mathbf{R} , respectively.

The Kalman filter is the minimum-variance state estimator. For systems with non-Gaussian noise, the Kalman filter is the minimum variance *linear* state estimator. The Kalman filter algorithms appear in the literature in several different forms and with varied notation, for example in references 7, 8, 9 and 10. A basic form of the filter will be presented here. Although in general the state transition matrix Φ and the observation matrix \mathbf{H} may vary with time, they will be shown as constants for this application.

The one-step predicted state estimate is

$$\hat{\mathbf{x}}^-(k+1) = \Phi\hat{\mathbf{x}}(k) + \bar{\mathbf{w}}(k)$$

with estimated error covariance

$$\mathbf{P}^-(k+1) = \Phi\mathbf{P}(k)\Phi^t + \mathbf{Q}(k)$$

where $\hat{\mathbf{x}}^-(k)$ is the predicted estimate of the state just *before* the current measurement is incorporated at time k , and $\mathbf{P}^-(k) \equiv E[\tilde{\mathbf{x}}^-(k)\tilde{\mathbf{x}}^-(k)^t]$, where $\tilde{\mathbf{x}}^-(k) \equiv \hat{\mathbf{x}}^-(k) - \mathbf{x}(k)$, $\mathbf{x}(k)$ being the true state vector. Similarly,

$\hat{\mathbf{x}}(k)$ is the updated estimate of the state just *after* the current measurement is incorporated at time k , and $\mathbf{P}(k) \equiv E[\tilde{\mathbf{x}}(k)\tilde{\mathbf{x}}(k)^t]$, where $\tilde{\mathbf{x}}(k) \equiv \hat{\mathbf{x}}(k) - \mathbf{x}(k)$, $\mathbf{x}(k)$ being the true state vector.

The measurements are then incorporated as follows. (The discrete time index k has been incremented at this point for convenience.) The Kalman gain, or blending factor, is given by

$$\mathbf{K}(k) = \mathbf{P}^-(k)\mathbf{H}^t[\mathbf{H}\mathbf{P}^-(k)\mathbf{H}^t + \mathbf{R}(k)]^{-1}$$

The state update is then

$$\hat{\mathbf{x}}(k) = \hat{\mathbf{x}}^-(k) + \mathbf{K}(k)[\mathbf{z}(k) - \mathbf{c}(k) - \mathbf{H}\hat{\mathbf{x}}^-(k)]$$

where $\mathbf{c}(k)$ represents the differential correction or the bias model. The true measurement-noise bias $\bar{\mathbf{v}}(k)$ differs from its estimate $\mathbf{c}(k)$ by the residual bias error $\bar{\mathbf{v}}_{res}(k) \equiv \bar{\mathbf{v}}(k) - \mathbf{c}(k)$. The error covariance update is given by

$$\mathbf{P}(k) = [\mathbf{I} - \mathbf{K}(k)\mathbf{H}]\mathbf{P}^-(k)$$

or, alternatively, reference 10 uses

$$\mathbf{P}(k) = \mathbf{P}^-(k) - \mathbf{K}(k)[\mathbf{H}\mathbf{P}^-(k)\mathbf{H}^t + \mathbf{R}(k)]\mathbf{K}(k)^t$$

or, for the stabilized Kalman filter, the expression

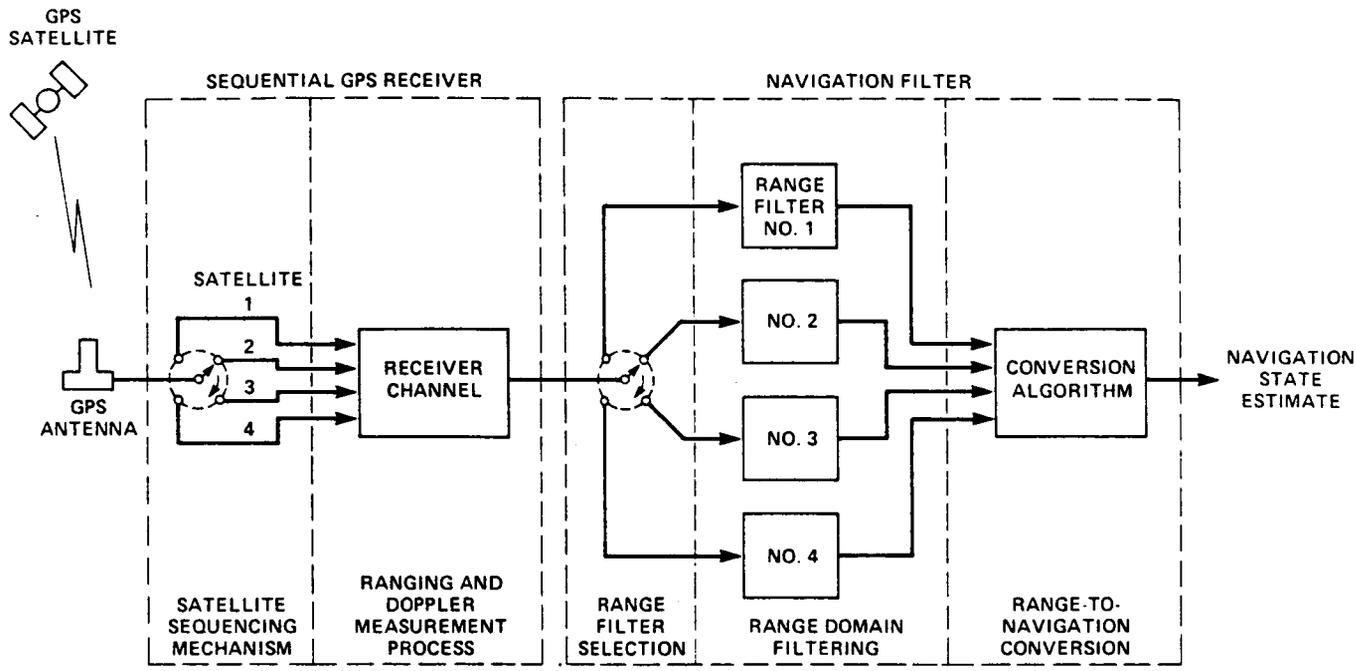
$$\mathbf{P}(k) = [\mathbf{I} - \mathbf{K}(k)\mathbf{H}]\mathbf{P}^-(k)[\mathbf{I} - \mathbf{K}(k)\mathbf{H}]^t + \mathbf{K}(k)\mathbf{R}(k)\mathbf{K}(k)^t$$

would be used, which reduces to the above expressions for $\mathbf{P}(k)$ when the Kalman gain is optimal. The stabilized Kalman filter might be used, for example, if the noise statistics are not well-known, because then the computed gain would not be truly optimal.

The updated state error covariance $\mathbf{P}(k)$ given by one of the preceding three equations is of course a symmetric matrix in theory, but will in practice become asymmetric because of finite-precision computation, which eventually causes filter divergence. This is avoided by forcing the symmetry of $\mathbf{P}(k)$ by averaging it with its transpose:

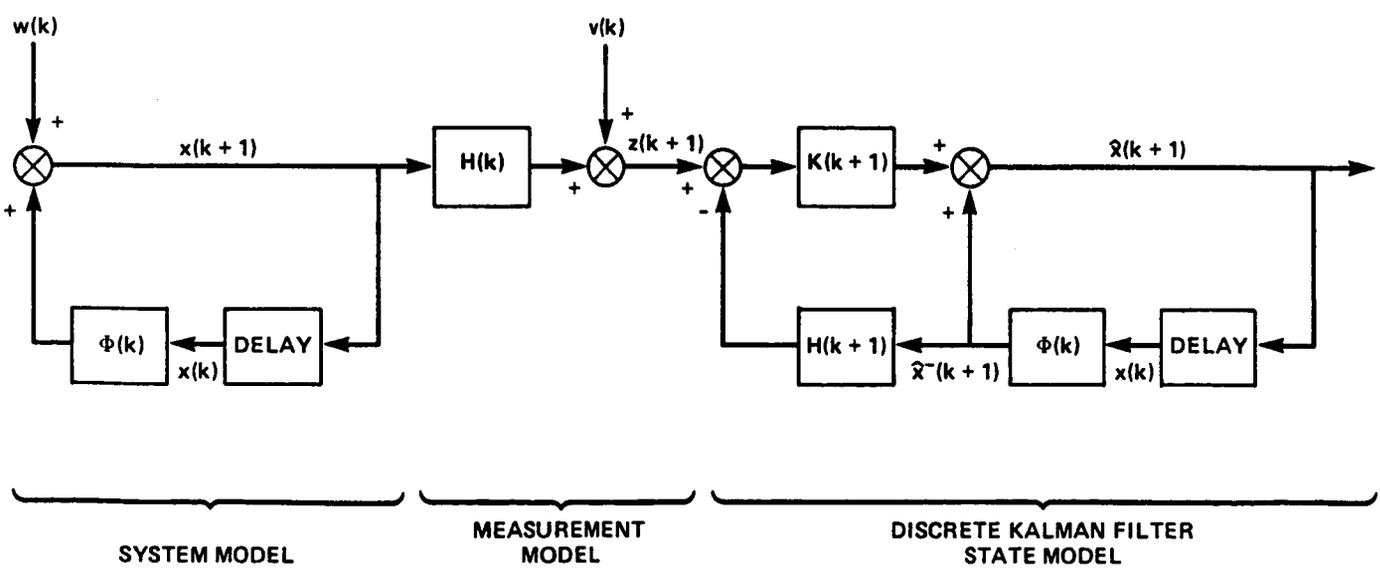
$$\mathbf{P}(k) := [\mathbf{P}(k) + \mathbf{P}(k)^t]/2$$

This completes the Kalman filter recursion. Figure 2 illustrates the structure of the physical model and of the Kalman filter state equations.



PAIELLI

Figure 1: Sequential GPS receiver with range filtering



PAIELLI

Figure 2: System model and Kalman filter state algorithm



Report Documentation Page

1. Report No. NASA TM-89418		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle Range Filtering for Sequential GPS Receivers with External Sensor Augmentation				5. Report Date April 1987	
				6. Performing Organization Code	
7. Author(s) Russell Paielli				8. Performing Organization Report No. A-87050	
				10. Work Unit No. 505-66-51-00	
9. Performing Organization Name and Address Ames Research Center Moffett Field, CA 94035				11. Contract or Grant No.	
				13. Type of Report and Period Covered Technical Memorandum	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546				14. Sponsoring Agency Code	
15. Supplementary Notes Point of Contact: Russell Paielli, Ames Research Center, M/S 210-9 Moffett Field, CA 94035 (415) 694-5454 or FTS 464-5454					
16. Abstract <p>The filtering of the satellite range and range-rate measurements from single channel sequential Global Positioning System receivers is usually done with an extended Kalman filter which has state variables defined in terms of an orthogonal navigation reference frame. An attractive suboptimal alternative is range-domain filtering, in which the individual satellite measurements are filtered separately before they are combined for the navigation solution. The main advantages of range-domain filtering are decreased processing and storage requirements and simplified tuning. Several range filter mechanization alternatives are presented, along with an innovative approach for combining the filtered range-domain quantities to determine the navigation state estimate. In addition, a method is outlined for incorporating measurements from auxiliary sensors such as altimeters into the navigation state estimation scheme similarly to the satellite measurements. A method is also described for incorporating inertial measurements into the navigation state estimator as a process driver.</p>					
17. Key Words (Suggested by Author(s)) GPS navigation, Kalman filter Kalman satellite navigation			18. Distribution Statement Unclassified-Unlimited Subject Category - 04		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of pages	22. Price